

REMARKS

Counsel for Applicants wishes to thank the Examiner for the courtesy of the interview of October 4, 2006, the substance of which is incorporated herein. For the reasons expressed in the interview, and set out below, it is respectfully submitted that the cited prior art to Latif et al. can neither anticipate nor render obvious the claimed invention, and favorable reconsideration of this application and withdrawal of the rejections are respectfully requested.

As pointed out during the interview, the present invention relates to storage methods for use in multi-protocol storage area networks (SANs). Conventional SANs comprise a network of a plurality of servers connected to a plurality of storage devices by storage switches. The servers, switches, storage devices and the interconnecting network usually all operate using different protocols, such as Fibre Channel ("FC"), iSCSI, and Ethernet (IP). Thus, it is necessary to translate between protocols to permit servers and storage devices to communicate if they (and the switches and network) do not operate on the same protocol. Translation between protocols requires significant memory and processor resources in conventional SANs, which causes delays in transmitting data and increases the cost, size and complexity of the system. The invention avoids these problems of conventional systems by performing protocol translation in an intelligent storage switch.

The Latif et al. Reference

The cited reference, U.S. Pat. No. 6,400,730 to Latif et al., illustrates a conventional SAN, and addresses the necessity of translating between different and incompatible protocols (see column 1, line 64 – column 2, line 12). Latif’s approach to protocol translation is, however, very different from the approach taken by the invention. As best described at column 2, line 55 – column 3, line 5, and shown in Figures 3 and 5, Latif converts an input packet at a source port having a first data or frame format (protocol), such as SCSI, Fibre Channel (“FC”) or Internet Protocol (“IP”), to an internal frame format which is used by his switch fabric for routing the packet through the switch, and then reconverts the packet at the switch output port back into the appropriate native format (SCSI, FC or IP) for the network or destination device connection to the destination port (see column 6, lines 44-57; and column 7, lines 59 - 61).

Latif teaches using a “Fibre Channel Protocol for SCSI” (“FCP”) as an internal switch format, converting FC, SCSI and IP data and commands into FCP frames, and then encapsulating the FCP frames in multiple IP packets (column 6, lines 24 – 25, and 56 – 65). This approach allows SCSI commands and data from a SCSI device to be transferred “directly using a parallel bus 106” to the FCP format and to the switch fabric (see column 6, lines 63 – 65, and Figure 3). Figures 6a – 6c illustrate FCP packet encapsulation into IP frames for transmission on an Ethernet interface.

The Claims Distinguish Over Latif et al.

This application contains independent Claims 1, 11, 12, 20, 34 and 49. Claims 1-18, 20-23, 34-36 and 49, comprising all independent claims, were rejected under 35 U.S.C. §102(e) as anticipated by Latif; and dependent Claims 19 and 24-33 were rejected under 35 U.S.C. §103(a) as obvious over Latif.

As pointed out during the interview, all independent claims, other than Claim 34, recite one of two limitations that are neither disclosed nor suggested by Latif, and that distinguish over Latif. These limitations are, as described during the interview and as will be explained more fully below, that the claimed method operates "without buffering" and that it operates "at wire speed". This is accomplished using an intelligent storage switch which affords storage methods that translate the data packets between protocols on-the-fly without buffering the data packets as is done in conventional approaches and at "wire speed". Accordingly, the invention minimizes the time required to process packets, and has a higher bandwidth and throughput rate than conventional systems. Independent Claim 34 distinguishes over Latif for other reasons that were pointed out during the interview, and which will be discussed below.

1. Without Buffering

Independent Claims 1, 12, 20 and 49, and dependent Claims 14 and 35 explicitly recite that all steps of the recited method are performed without buffering. (Claim 49 recites that the instructions do not require the packet to be buffered.) As

pointed out during the interview, Latif does not disclose or suggest that his method operates “without buffering”. In fact, Latif teaches just the opposite, i.e., using a buffer.

In performing his protocol translating method, Latif discloses encapsulating packets into an FCP data frame having a maximum size (column 9, lines 28-30) and teaches that his switch 235 uses a Buffer to Buffer Receive Data Field sized as is necessary to force communication with data frames that fit within an IP packet carried on an Ethernet. (See, column 9, lines 48 – 57, and column 10, lines 18 – 42. See also, column 15, lines 14-16.) Thus, Latif clearly contemplates and teaches buffering of data packets as part of his translation process between different protocols.

The reference in Latif, at column 11, lines 55-63 to processing a packet “directly”, referred to by the Office in its Office Action, does not mean that packets are processed “without buffering”. Rather, Latif is merely describing at that location that if the destination device of an FCP packet is an SoIP (“Storage Over Internet Protocol”) device, the IP encapsulated FCP packet may be processed “directly”, i.e., “de-encapsulating and processed as an FCP packet” by the destination device. Latif’s disclosure of “processing a packet directly” has no relation to, and discloses or suggests nothing about, processing a packet “without buffering”. In fact, Latif expressly teaches buffering at column 15, lines 14-16, where he discloses that it is “necessary to buffer the entire frame to determine the length and checksum and write them into the header”.

Accordingly, Latif cannot anticipate or render obvious any of independent Claims 1, 12, 20 or 49, or dependent Claims 2 – 11, 13 - 18, 21 – 22 or 35, and reconsideration of the rejection of these claims is respectfully requested.

2. At Wire Speed

The second limitation that Latif does not disclose (or suggest), and which is recited in independent Claim 11, is that the processing occurs at “wire speed”.

The specification defines the claim term “wire speed” to mean that packets are processed and translated between the different protocols without introducing any more latency than would be introduced by merely switching or routing functions (see specification, page 6, paragraph [0015], lines 15-22; and page 13, paragraph [0062], lines 12-15). In particular, the specification (see page 13, paragraph [0062], lines 15-16), defines the term “wire speed” to be a speed measured by the connection to a given port of the switch, and that “wire speed” means that the switch of the invention can translate packets at the maximum packet throughput rate for the type of connection to a switch port. Thus, wire speed is defined relative to the speed of the connection to a switch port, and requires processing and translating a packet at the maximum packet throughput rate at the switch port. In order to process packets at wire speed, the invention will process packets “without buffering”, and will process the packets on-the-fly (see specification, page 18, paragraph [0078]; page 33, paragraph [0127]).

The specification elaborates on the definition of wire speed with examples. As described on page 13, if the connection to a port is an OC-48 connection (2.5 bits per nsec), wire speed means that translation must keep up with that data rate. Similarly, if a switch port is a 1 Gb Ethernet, wire speed processing will be 1 bit per nsec; and if the port is 2 Gb Fibre Channel the wire speed will be 5 μ sec per Kilobyte.

Latif does not disclose, either explicitly or implicitly, or suggest that packet translation and processing be performed at the speed of the connection to the input port of the received packet, i.e., wire speed, as claimed. In fact, as pointed out above, Latif contemplates buffering of packets during reformatting, and, accordingly, cannot process packets at wire speed. Independent Claim 11 and dependent Claims 2, 13, 21 and 36 all explicitly recite this limitation. Accordingly, Latif also cannot anticipate, or render obvious, these claims or the claims dependent thereon, and reconsideration of the rejection of Independent Claim 11 and dependent Claims 2, 13, 21 and 36 is respectfully requested.

Claim 34

Independent Claim 34 distinguishes over Latif for other reasons. This claim calls for a method for use in a storage system wherein an ingress linecard receives information about a virtual target address of a packet, including a flowID, puts a virtual target descriptor (VTD) identifier and the flowID into a header of the packet, and routes the packet to an egress line card in accordance with the flowID. The egress linecard uses the virtual target descriptor identifier to identify information about the physical

target to convert a virtual target block address to a physical target block address, and to translate, if necessary, the format of the packet from a first protocol to a second protocol.

Latif does not disclose the use of a virtual target descriptor identifier and flowID, as claimed. At col. 2, line 66 – col. 3, line 5, referenced by the Office, Latif merely discloses converting a received packet to an internal format (as described above), routing the packet through the switch fabric to a destination port, and reconvertng the internal packet format back to the native format of a device connected to the port.

As pointed out during the interview, Figures 6a – 6c of Latif illustrate the mapping of fields to perform FCP packet encapsulation, and Figure 6c (referred to by the Office in its rejection) merely illustrates the fields for the source port length and the destination port checksum (see column 8, lines 38-61). Latif does not disclose or suggest use of a virtual target descriptor identifier and flowID, as claimed. Accordingly, Latif cannot anticipate or render obvious Claim 34, and reconsideration of the rejection of Claim 34 is respectfully requested.

In view of the foregoing, it is respectfully submitted that the Latif cannot anticipate or render obvious any of Claims 1-36 and 49, and that these claims are allowable. Accordingly, reconsideration and withdrawal of the rejections of the claims under 35 U.S.C. §§ 102 and 103 are requested .

In view of the foregoing, it is respectfully submitted that this application is in condition for allowance and early allowance of all claims is solicited.

If any issues remain, it is requested that the Examiner telephone the undersigned prior to issuing an Office Action.

Dated: October 6, 2006

Respectfully Submitted,

/Barry N. Young/

Barry N. Young
Attorney for Assignee
Reg. No. 27,744

Customer No. 48789

Law Offices of Barry N. Young
Court House Plaza, Suite 410
260 Sheridan Avenue
Palo Alto, CA 94306-2047
Phone: (650) 326-2701
Fax: (650) 326-2799
byoung@young-iplaw.com